arXiv:1206.3461v1 [astro-ph.IM] 15 Jun 2012

Avoiding selection bias : A unified treatment of thresholded data

C. Messenger*
School of Physics and Astronomy, Cardiff University, Queen's Buildings, The Parade, CF24 3AA

J. Veitch†
Nikhef, Science Park 105, Amsterdam 1098 XG, Netherlands

(Dated: Fri Jun 15 13:02:59 2012 +0100)
(commitID: dcd822c4276f5b7492f102369a7795994f2f2993)
(LIGO-P1200062-v2)

When searching for populations of rare and/or weak signals in noisy data, it is common to use a detection threshold to remove marginal events which are unlikely to be the signals of interest; or a detector might have limited sensitivity, causing it to not detect some of the population. In both cases a selection of data has occurred, which can potentially bias any inferences drawn from the remaining data, and this effect must be corrected for. We show how the selection bias is naturally avoided by using the full information from the search, considering both the selected data and our ignorance of the data that are thrown away, and considering all relevant signal and noise models. This approach produces unbiased estimates of parameters even in the presence of false alarms and incomplete data.

Introduction— There are many areas of the physical sciences where experiments or observations are performed in order to measure or search for populations of events or signals. For example, in astrophysics there are searches for Type Ia supernovae (SNIe) [1], gravitational-waves [2], extra-solar planets [3] and high energy neutrinos [4, 5]; and galaxy cluster surveys [6, 7]. Often the aim is to use the resulting set of observations to infer global parameters of interest, e.g. quantities that characterise the population distribution. For the case of SNIe for example, one may aim to infer cosmological parameters from luminosity distance and redshift measurements. In this paper we focus on these global quantities rather than properties of individual events and show how to derive an unbiased inference from data sets that are subject to *selection bias*.

Many experiments are subject to selection bias, which we define as *the influence that a restriction on the measured data has on quantities inferred from a population of events or signals*. Data could be selected either deliberately by the experimenter, in order to reduce the background of noise events, or inadvertently through limitations of the detection equipment, such as a saturation point or limited sensitivity to weak signals. An example in the astrophysical literature is a flux-limited survey [8]. We draw a careful distinction between selection bias and biases induced by an incorrect application of prior information [9, 10], as the terms "Malmquist bias" [11–13] or "Lutz-Kelker bias" [14] are sometimes used to describe both types of bias.

The problem of selection bias is present in many areas and in some fields is extensively documented. In astronomy, Kelly [15] takes a statistical approach to the bias problem which is closest to ours, considering both detected and non-detected events. In cosmology, the traditional approach has been to produce unbiased *estimators* of individual event parameters [16–20] which are then used as input for inference of global quantities. Thresholding can be applied to reduce the number of false alarms from a background population to a tolerable level, for example thresholding on radio pulse width is applied to data within the LOFAR NuMoon [21] project. The ANTARES [22] neutrino experiment applies a threshold on a fitting parameter to optimise the detection efficiency of cosmic neutrinos whilst minimising the effect of the atmospheric neutrino and muon background. In gravitational wave searches [23], thresholds are set on the signal-to-noise-ratio (SNR) of an event to reduce the volume of data to be processed, and additional cuts and vetoes are applied within the data-analysis pipeline in order to minimise the false alarm rate due to detector noise.

In this letter we apply Bayesian methods to provide a prescription for completely avoiding selection bias. It has been previously shown [15], for astronomical data-analysis, that to avoid bias in a signal dominated environment (i.e. no false alarms or no events due to noise alone) one should account for both the detections *and* the false dismissals within an experiment. We generalise and complete this approach by also considering the probabilities of detection and dismissal for all other relevant models, most importantly the noise model. Our approach correctly includes the information content of the data which are thrown away, making no assumption that they contain only noise, nor do we assume that detected events are produced only by true signals.

We will first derive general expressions for the posterior probability density functions (pdfs) on the parameters of interest for any threshold-influenced analysis. We follow this with an example application to a simulated case of cosmological inference using standard candle transients. We conclude with a discussion of the uses and implications of this approach.

General Formulation— We consider an experiment in which many measurements are made, each producing data which may contain potential detections, and which are then compared with a detection threshold. We refer to any measurements which pass the threshold as "triggers", since in most practical cases we will include one or more noise-only models alongside the true signal model. Any measurements which do not pass the threshold are discarded. We will derive an expression for calculating global parameters based on the triggers, but which also takes into account the cases when a trigger was not produced. These non-trigger cases also carry information about the source population, and one cannot assume the lack of a trigger implies the lack of true signal, or vice-versa.

We divide the experiment into multiple measurements, whose results we suppose to be statistically independent of one another. We will denote the outcome of the experiment D. The threshold is set on the measured quantity x, which we will call the "detection statistic". In general x may be a maximum-likelihood statistic or simply a direct measurement of a physical parameter value but in all circumstances this measurement will be subject to noise. For simplicity of notation we only consider situations in which values of x are restricted to be greater than a threshold $x_{\rm th}$ [24]

If the statistic passes a threshold $x > x_{\rm th}$ a trigger is produced, and the value of x is recorded; we denote this case D^+, and $D = (D^+, x)$. If on the other hand $x \leq x_{\rm th}$, no trigger is produced and no value for x is recorded and $D = D^-$. We define $\{\mathcal{H}\}$ as our set of plausible models and must include at least two models with one describing the presence and the other the absence of a signal. In general the method extends to many alternative signal models. When specifying a set of models in this way we allow that any model may be true, and that the trigger has a non-zero probability under all of the models. For any single model $\mathcal{H}$ the likelihood function defining the distribution of the statistic x is expressed as $p(x|\mathcal{H}, \boldsymbol{\theta}_{\mathcal{H}}, \boldsymbol{\gamma}, I)$ where $\boldsymbol{\theta}_{\mathcal{H}}$ are the parameters unique to this model, and $\boldsymbol{\gamma}$ are the *global* parameters which are common to all trials (following standard Bayesian notation we use I throughout to represent additional information regarding the experiment, and the global parameters $\boldsymbol{\gamma}$). Upon performing the single trial, we are interested in inferring the parameters $\boldsymbol{\gamma}$. In the case of a statistic above threshold, D^+, we can write the posterior distribution on $\boldsymbol{\gamma}$ as

$$\begin{aligned} p(\boldsymbol{\gamma}|D^+, x, I) &= \sum_{\{\mathcal{H}\}} p(\boldsymbol{\gamma}, \mathcal{H}|D^+, x, I) \\ &= \frac{p(\boldsymbol{\gamma}|I)}{P(x|I)} \overbrace{\sum_{\{\mathcal{H}\}} p(x|\boldsymbol{\gamma}, \mathcal{H}, I) P(\mathcal{H}|\boldsymbol{\gamma}, I)}^{\text{trigger likelihood } p(D^+, x|\boldsymbol{\gamma}, I)} \end{aligned} \tag{1}$$

where $P(\mathcal{H}|\boldsymbol{\gamma}, I)$ represents the prior probability on the model $\mathcal{H}$ conditional on the global parameters. Here we are summing the trigger likelihoods over the possible models, and we do not assume that the model priors are independent of the global parameters, which will be important when discussing signal abundance and rates. We have used the fact that the probability of a trigger conditional on the value x is strictly equal to $H(x - x_{\rm th})$ (where $H(x)$ is the Heaviside step function) which is equal to unity for $x > x_{\rm th}$. If the model has local parameters $\boldsymbol{\theta}_{\mathcal{H}}$, we calculate the factor $p(x|\boldsymbol{\gamma}, \mathcal{H}, I)$, i.e. the likelihood of the data conditional on the model and the global parameters, by marginalisation:

$$p(x|\boldsymbol{\gamma}, \mathcal{H}, I) = \int_{\Theta_{\mathcal{H}}} d\boldsymbol{\theta}_{\mathcal{H}} p(x|\boldsymbol{\theta}_{\mathcal{H}}, \boldsymbol{\gamma}, \mathcal{H}, I) p(\boldsymbol{\theta}_{\mathcal{H}}|\boldsymbol{\gamma}, \mathcal{H}, I), \tag{2}$$

where $\Theta_{\mathcal{H}}$ is the local parameter space associated with the local parameter $\theta_{\mathcal{H}}$ and the model $\mathcal{H}$. Note that we do not modify this likelihood function to reflect our knowledge of the threshold.

In the alternative case D^- where we discarded the specific value of x, we do still know that $x \leq x_{\rm th}$. For a non-trigger we can therefore write the posterior probability on $\boldsymbol{\gamma}$ (analogous to Eq. (1)) as

$$p(\boldsymbol{\gamma}|D^-, I) = \frac{p(\boldsymbol{\gamma}|I)}{P(D^-|I)} \overbrace{\sum_{\{\mathcal{H}\}} P(D^-|\boldsymbol{\gamma}, \mathcal{H}, I) P(\mathcal{H}|\boldsymbol{\gamma}, I)}^{\text{non-trigger likelihood } p(D^-|\boldsymbol{\gamma}, I)} \tag{3}$$

As the value of x is unknown the posterior on $\boldsymbol{\gamma}$ must be marginalised over both the local parameters $\boldsymbol{\theta}_{\mathcal{H}}$ and the unknown data x in the range $x \leq x_{\rm th}$. Therefore only the information D^- appears in the likelihood of a non-trigger conditional on the model and the global parameters:

$$P(D^-|\boldsymbol{\gamma}, \mathcal{H}, I) = \int_{\Theta_{\mathcal{H}}} d\boldsymbol{\theta}_{\mathcal{H}} \int_{x \leq x_{\rm th}} dx\, p(x|\boldsymbol{\theta}_{\mathcal{H}}, \boldsymbol{\gamma}, \mathcal{H}, I) p(\boldsymbol{\theta}_{\mathcal{H}}|\boldsymbol{\gamma}, \mathcal{H}, I) \tag{4}$$

The quantity derived in Eq. (4) is related to what is known as the *selection function* or *detection efficiency*, defined as the probability of a trigger given the true signal parameters for a given signal model. Here it appears in a modified form: the probability of a *non-trigger* given a set of global parameters and a model. It is only in this term that the detection threshold $x_{\rm th}$ actually appears. In practice this quantity (just like the detection efficiency) can be difficult to estimate in some cases.

An ensemble of trials— Using Eqns. (1) and (3) together with Bayes' theorem we are able to isolate the single trial likelihood functions (identified by the overbraces) in both the trigger and non-trigger cases, $P(D^+, x|\boldsymbol{\gamma}, I)$ and $p(D^-|\boldsymbol{\gamma}, I)$ respectively. We can now also consider analysing an ensemble of N trials denoted as $\{D\}$, where we index each D_i using $i \in (1, N)$. The posterior probability for the global parameters after N independent trials are performed is simply

$$p(\boldsymbol{\gamma}|\{D\}, I) = \frac{p(\boldsymbol{\gamma}|I)}{P(\{D\}|I)} \prod_{i=1}^{N} p(D_i|\boldsymbol{\gamma}, I). \tag{5}$$

We can further decompose this expression into trigger D^+ and non-trigger D^- factors,

$$p(\boldsymbol{\gamma}|\{D\}, I) = \frac{p(\boldsymbol{\gamma}|I)}{P(\{D\}|I)} \prod_{j=1}^{n} p(D^+, x_j|\boldsymbol{\gamma}, I_j) \prod_{k=1}^{N-n} P(D^-|\boldsymbol{\gamma}, I_k), \tag{6}$$

where j and k index the n triggers and the $N-n$ non-triggers respectively, and I_j, I_k allow for different background information (such as time-varying detector sensitivity) in each measurement.

We now focus on the common case where there are only two possible models: the presence ($\mathcal{H}^+$) or absence ($\mathcal{H}^-$) of a single signal per trial. Equation (6) becomes

$$\begin{aligned} p(\boldsymbol{\gamma}|\{D\},I) = & \frac{p(\boldsymbol{\gamma}|I)}{P(\{D\}|I)} \\ & \times\prod_{j=1}^{n}\Big[p(D^+,x_j|\boldsymbol{\gamma},\mathcal{H}^+,I_j)P(\mathcal{H}^+|\boldsymbol{\gamma},I_j) \\ & \quad +p(D^+,x_j|\boldsymbol{\gamma},\mathcal{H}^-,I_j)P(\mathcal{H}^-|\boldsymbol{\gamma},I_j)\Big] \\ & \times\prod_{k=1}^{N-n}\left[P(D^-|\boldsymbol{\gamma},\mathcal{H}^+,I_k)P(\mathcal{H}^+|\boldsymbol{\gamma},I_k)\right. \\ & \quad \left.+P(D^-|\boldsymbol{\gamma},\mathcal{H}^-,I_k)P(\mathcal{H}^-|\boldsymbol{\gamma},I_k)\right]. \end{aligned} \tag{7}$$

To relate this to more standard statistical terms, consider the key components which make up the expression, each of which are in reference to single trials:

- $p(D^+,x_j|\boldsymbol{\gamma},\mathcal{H}^+,I_j)$ is the likelihood of producing a trigger with statistic value x_j when a signal is truly present.
- $p(D^+,x_j|\boldsymbol{\gamma},\mathcal{H}^-,I_j)$ is the likelihood of producing a trigger with statistic value x_j from only the background distribution.
- $P(D^-|\boldsymbol{\gamma},\mathcal{H}^+,I_k)$ is the likelihood that a true signal does not produce a trigger, i.e. the *false dismissal likelihood* (or 1−detection efficiency).
- $P(D^-|\boldsymbol{\gamma},\mathcal{H}^-,I_k)$ is the likelihood that no trigger is produced when no signal is present, i.e. the *true dismissal likelihood.*

The abundance or true rate of signal events is important in the two model case, when the two models are signal and background. As the two models give a complete description of the data, we can write $P(\mathcal{H}^+|\boldsymbol{\gamma},I) = 1 - P(\mathcal{H}^-|\boldsymbol{\gamma},I)$. The actual rate of signals, or a quantity from which it can be derived, must be included as one of the global parameters $\boldsymbol{\gamma}$, and if it is not known then it can be estimated from the ensemble of trials. This makes the prior probabilities for the models dependent on the global parameters $\boldsymbol{\gamma}$, so the joint posterior distribution for the rate will not in general be separable from the rest of global parameters. Alternatively, if the true rate of signals is known, it can be substituted into these expressions instead. Neglecting this effect is a potential source of bias in inference of both the rate and other global parameters. We will show a specific example of this calculation, including the rate estimate below.

The key idea of our approach (in terms best suited to the two model problem) is to use the true and false dismissal probabilities to incorporate the information from the absence of triggers. This requires that these quantities be either calculated or measured. Without these, any observations which can produce non-triggers risk being either biased or sub-optimal and may require ad-hoc corrections. It is only by acknowledging and accounting for ignorance that we can derive the correct results. Equation (7) contains all the parts necessary to infer global parameters without bias, making full use of the information that is kept, and accounting for that which is deliberately (or necessarily) discarded.

A standard candle example— To illustrate our method, we will consider a toy model in which a telescope performs a flux-limited survey for transient standard-candle events, much like SNIe surveys. Although our model is simple and lacks the detail of real observations, it will serve to show the salient points of the method. In the model we are interested in inferring the global parameters R, the astrophysical event rate per unit co-moving volume, H_0, the Hubble constant, and Ω_m, the matter density, from simulated trials. We assume that the telescope is able to scan a patch of the sky $\alpha \equiv 2\times 2\ \mathrm{deg}^2$ in $\Delta t = 60$ seconds and that this constitutes a single trial. For each trial, the "data" comprises a measurement of the distance modulus μ and the redshift z of any event that passes the thresholds. In this case we apply thresholds on satisfying $\mu \leq \mu^{\mathrm{th}}$ and $z \leq z^{\mathrm{th}}$ where $\mu^{\mathrm{th}} = 45$ is the distance modulus threshold and $z^{\mathrm{th}} = 2$ is a redshift threshold. Our measurements of μ and z are Gaussian random variables with standard deviations $\sigma_\mu = 0.5$ and $\sigma_z = 0.02z$ respectively, i.e. the fixed distance modulus uncertainty which can be interpreted as either an intrinsic scatter in the standard candle luminosity or as a measurement error, and a redshift dependent z uncertainty. We suppose that events arise through either a signal model $\mathcal{H}^+$ or through a noise model $\mathcal{H}^-$. According to the signal model we define the distance modulus of a source as $\mu = 5\log_{10}(d_{\mathrm{L}}) + 25$ where the luminosity distance d_{L} is measured in Mpc and defined using $d_{\mathrm{L}} = (1+z)d_{\mathrm{M}}$ with the co-moving radial distance $d_{\mathrm{M}} = c\int_0^z H(z')^{-1}dz'$ with c as the speed of light. The Hubble parameter $H(z)$ is given by $H(z) = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_k(1+z)^2 + \Omega_\Lambda}$ where for this example we assume a flat universe $\Omega_k = 0$ and fix dark energy parameters $w_0 = -1, w_a = 0$ such that $\Omega_\Lambda = 1 - \Omega_m$. The model includes only one local parameter, z_0, the true redshift of each signal. Under the noise model, we assume that events are drawn from a unit variance Gaussian distribution with mean distance modulus $\mu_{\mathcal{H}^-} = 46.8$ and can be interpreted as arising from a redshift-independent detector noise in the telescope at the telescope flux limit. All data recorded above threshold due to the noise model will have a corresponding redshift value drawn from the redshift prior since we assume an existing host galaxy would be incorrectly identified as the host.

We choose a maximum redshift $z^{\max} = 4$ such that the cosmological volume enclosed is large enough that with a high confidence the observed comoving volume $V = \alpha\int_0^{z^{\max}} d_{\mathrm{M}}^2(z')/[(1+z')H(z')]\,dz'$ contains all detectable signals, i.e. those true signals that we would expect to pass the thresholds. Within this sphere we assume a uniform distribution of sources in co-moving volume giving a prior distribution on the observed redshift of any given source as $p(z|\boldsymbol{\gamma},I) = \kappa d_{\mathrm{M}}^2(z)/(1+z)H(z)$ where κ is a constant of proportionality ensuring the correct normalisation of the prior

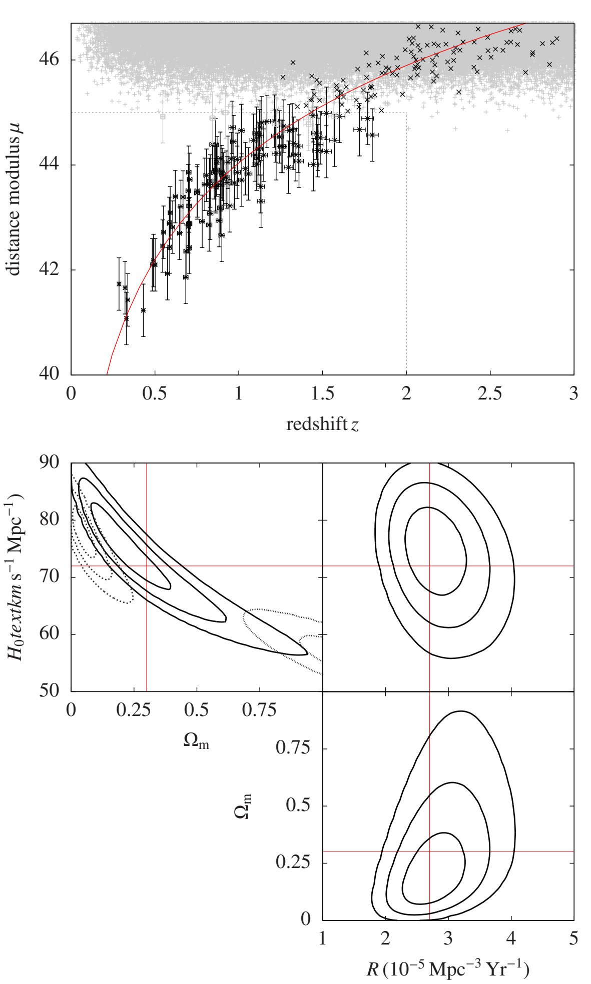

FIG. 1. (upper) The simulated data set consisting of 115 potential detections or triggers (those with errorbars) with magnitudes less than $\mu^{\mathrm{th}} = 45$ and redshifts less than $z^{\mathrm{th}} = 2$ (bounded by grey dashed lines) from a total of 131400 independent trials. Of those, 108 were signals and 7 were noise events. Those that did not pass the threshold are shown without error bars. For both triggers and non-triggers we further subdivide measurements into signal model $\mathcal{H}^+$ (black) and noise model $\mathcal{H}^-$ (grey). The red curve represents the true relation between the redshift and distance modulus. (lower) Joint probability contours on the combined global parameters of the cosmological parameters H_0 and Ω_m and the astrophysical rate R. The true values of these parameters are indicated by the solid red lines. The contours represent regions of equal enclosed probability at values corresponding to 1,2 and 3 standard deviations. The black contours correspond to the application of our method accounting for both triggers and non-triggers and both the signal and noise models. The dashed grey contours represent results from an analysis accounting for triggers-only and signal model only. The dotted grey contours represent an analysis accounting for triggers-only but including both the signal and noise models.

$\int_0^{z^{\max}} dz\; p(z|\boldsymbol{\gamma}, I) = 1$. We note that this prior is dependent on the global parameter Ω_m and identical for both models. We simulate 1 year of trials assuming 360 observations per night with a true astrophysical event rate (per comoving volume per unit observed time) of $R = 2.7 \times 10^{-5}\,\mathrm{Mpc}^{-3}\mathrm{yr}^{-1}$ with cosmological parameters $H_0 = 72\,\mathrm{km\,s}^{-1}\mathrm{Mpc}^{-1}$ and $\Omega_m = 0.3$. The triggers (in addition to non-triggers) are shown in the upper plot of Fig. 1 where we reveal which are produced in the presence and absence of a true signal.

First consider the probability of obtaining a non-trigger per trial under the assumption of the signal model $\mathcal{H}^+$,

$$P(D^-|\boldsymbol{\gamma},\mathcal{H}^+,I) = 1 - \int_0^{\mu^{\mathrm{th}}} d\mu \int_0^{z^{\mathrm{th}}} dz \int_0^{z^{\max}} dz_0\; p(\mu,z,z_0|\boldsymbol{\gamma},\mathcal{H}^+,I)$$
$$= 1 - \int_0^{z^{\max}} dz_0\, \frac{p(z_0|\boldsymbol{\gamma},I)}{4} \mathrm{erfc}\left(\frac{\mu(z_0,\boldsymbol{\gamma}) - \mu^{\mathrm{th}}}{\sqrt{2\sigma_\mu^2}}\right) \mathrm{erfc}\left(\frac{z_0 - z^{\mathrm{th}}}{\sqrt{2\sigma_z^2(z_0)}}\right). \tag{8}$$

The complementary quantity, the probability density associated with detection for the signal model, for the j'th trial is

$$p(\mu_j,z_j|\boldsymbol{\gamma},\mathcal{H}^+,I) = 1 - \int_0^{z^{\max}} dz_0\; p(\mu_j,z_j,z_0|\boldsymbol{\gamma},\mathcal{H}^+,I)$$
$$= \int_0^{z^{\max}} dz_0\, \frac{p(z_0|\boldsymbol{\gamma},I)}{2\pi\sigma_\mu\sigma_z(z_0)} \exp\left(-\frac{\left(\mu_j - \mu(z_0)\right)^2}{2\sigma_\mu^2} - \frac{\left(z_j - z_0\right)^2}{2\sigma_z^2(z_0)}\right). \tag{9}$$

For the noise model we have similar expressions but in this case the distance modulus is independent of the redshift. For non-trigger and trigger cases under the noise model we have

$$P(D^-|\boldsymbol{\gamma},\mathcal{H}^-,I) = 1 - \int_0^{\mu^{\mathrm{th}}} d\mu \int_0^{z^{\mathrm{th}}} dz \int_0^{z^{\max}} dz_0\; p(\mu,z,z_0|\boldsymbol{\gamma},\mathcal{H}^-,I)$$
$$= 1 - \frac{1}{4}\mathrm{erfc}\left(\frac{\mu_{\mathcal{H}^-} - \mu^{\mathrm{th}}}{\sqrt{2\sigma_\mu^2}}\right) \int_0^{z^{\max}} dz_0\, p(z_0|\boldsymbol{\gamma},I)\, \mathrm{erfc}\left(\frac{z_0 - z^{\mathrm{th}}}{\sqrt{2\sigma_z^2(z_0)}}\right) \tag{10}$$

$$p(\mu_j,z_j|\boldsymbol{\gamma},\mathcal{H}^-,I) = \int_0^{z^{\max}} dz_0\; p(\mu_j,z_j,z_0|\boldsymbol{\gamma},\mathcal{H}^-,I)$$
$$= \frac{e^{-(\mu_j - \mu_{\mathcal{H}^-})^2/2\sigma_\mu^2}}{2\pi\sigma_\mu} \int_0^{z^{\max}} dz_0 \exp\left(-\frac{(z_j - z_0)^2}{2\sigma_z^2(z_0)}\right) \frac{p(z_0|\boldsymbol{\gamma},I)}{\sigma_z(z_0)}. \tag{11}$$

We define the model priors assuming the following limit. For observation lengths Δt the probability of there being more than a single true signal event (above threshold *or* otherwise) within the observable volume V is $\ll 1$. Since the occurrence of a signal event during a trial is governed by the Poisson distribution, the prior probability of the signal model is accurately approximated as $P(\mathcal{H}^{+}|\gamma, I) = RV\Delta t$. We consider our two models as exhaustive and as such the complementary prior noise model probability is given by $P(\mathcal{H}^{-}|\gamma, I) = 1 - RV\Delta t$.

As shown in the lower plots of Fig. 1 we are able to simultaneously infer a joint posterior probability on the global parameters H_0, Ω_m and R. The model can be made arbitrarily more complicated by including any number of other unknown parameters, global or local, e.g. the intrinsic scatter in distance modulus σ_μ or additional cosmological parameters.

Two additional results are included with the aim of highlighting the biases that can occur when failing to incorporate all relevant information into an analysis. The first of these results corresponds to the naive case where only the triggers are considered and only the signal model is assumed. These results are calculated by evaluating Eq. 9 for each trigger and therefore ignores the information contained within both the model prior and the non-triggers. In this case the subsequent cosmological parameter estimates are highly bias towards small H_0 and large Ω_m values with R not estimated as it is present only in the model priors. This dominant cause of this bias is the truncation of the μ values by the threshold leading to a disproportionate number of triggers at higher redshift. Larger values of Ω_m lead to an increase in the downward curvature of the redshift-distance modulus relation better matching the data at higher redshifts whilst smaller H_0 values oppose this effect at low redshifts.

The second bias result corresponds to an analysis that considers only triggers but assumes both a signal and noise model. The resulting cosmological parameter estimates now incorporate the information present within the model priors but ignore the information in the non-triggers and subsequently yield a significant under-estimate of Ω_m and over-estimate of R (the astrophysical rate estimate has become completely saturated at the upper prior boundary and is not shown in Fig. 1). In these examples the non-trigger contribution to the result is dominated by the likelihood of obtaining a non-trigger from the noise model. For this case the missing component of the likelihood would primarilly act to shift estimates of Ω_m and R towards larger and smaller values respectively.

Discussion— Using a simple application of Bayesian inference we are able to derive a general formalism for analysing data which have been subjected to a threshold, either from the nature of an instrument or artificial by design. The resulting expressions account for both the knowledge of and the ignorance of the actual result of an experiment. The method succeeds for any number of false alarms and reduces to known formulae in limiting cases such as high signal to noise ratio, or zero false alarm rate. We find that knowledge of the true rate or abundance of signals represents crucial information when making inferences with an ensemble of trials, and that it can therefore be inferred from them. The simplicity and consistency of the Bayesian approach requires no ad-hoc corrections to be applied, and each term follows naturally from careful consideration of the problem. Prior knowledge of signal and noise models must be incorporated, but measured values of efficiency and false alarm rates can be used when they are difficult to derive analytically. Our method has applications in many noise-limited fields of research, ranging from particle physics to cosmology. It also allows one to study the change in precision of estimates when the threshold is lowered or raised, while retaining an unbiased result, and this could be useful when determining the optimal tradeoff between sensitivity and amount of data to retain.

Acknowledgments—The authors are grateful to J. Clark, T. Dent, S. Fairhurst, I. Harrison, M. Hendry, D. Keitel, T. G. F. Li, W. Del Pozzo, G. Pratten, R. Prix, C. Röver, B. S. Sathyaprakash, P. Sutton, M. Vallisneri C. Van Den Broeck, and G. Woan for useful discussions and comments.

---

* chris.messenger@astro.cf.ac.uk
† johnv@nikhef.nl

[1] M. Hicken, W. M. Wood-Vasey, S. Blondin, P. Challis, S. Jha, P. L. Kelly, A. Rest, and R. P. Kirshner, The Astrophysical Journal, **700**, 1097 (2009).
[2] B. Abbott *et al.*, Phys. Rev. D, **77**, 062002 (2008).
[3] S. Udry and N. C. Santos, Annual Review of Astronomy & Astrophysics, **45**, 397 (2007).
[4] M. Ageron *et al.*, Nuclear Instruments and Methods in Physics Research A, **656**, 11 (2011).
[5] R. Abbasi *et al.*, arXiv.org, **astro-ph.IM** (2010).
[6] S. W. Allen, A. E. Evrard, and A. B. Mantz, Annual Review of Astronomy & Astrophysics, **49**, 409 (2011).
[7] C. Cunha, Physical Review D, **79**, 63009 (2009).
[8] T. J. Loredo, Statistical Challenges in Modern Astronomy IV ASP Conference Series, **371**, 121 (2007).
[9] M. A. Hendry and J. F. L. Simmons, Astronomy and Astrophysics (ISSN 0004-6361), **237**, 275 (1990).
[10] T. J. Loredo and M. A. Hendry, "Bayesian multilevel modelling of cosmological populations," in *Bayesian Methods in Cosmology*, edited by Hobson, M. P., Jaffe, A. H., Liddle, A. R., Mukeherjee, P., & Parkinson, D. (Cambridge University Press, 2010) p. 245.
[11] A. S. Eddington, Monthly Notices of the Royal Astronomical Society, **73**, 359 (1913).
[12] K. G. Malmquist, Medd. Lund. Astron. Obs., **20** (1920).
[13] A. S. S. Eddington, Monthly Notices of the Royal Astronomical Society, **100**, 354 (1940).
[14] T. E. Lutz and D. H. Kelker, Publications of the Astronomical Society of the Pacific, **85**, 573 (1973).
[15] B. C. Kelly, Astrophysical Journal, **665**, 1489 (2007.).
[16] S. D. Landy and A. S. Szalay, Astrophysical Journal, **391**, 494 (1992).
[17] A. Mantz, S. W. Allen, D. Rapetti, and H. Ebeling, Monthly Notices of the Royal Astronomical Society, **406**, 1759 (2010).
[18] M. A. Hendry and J. F. L. Simmons, Astrophysical Journal, **435**, 515 (1994).
[19] J. A. Willick and M. A. Strauss, The Astrophysical Journal, **507**, 64 (1998).
[20] D. Lynden-Bell, S. M. Faber, D. Burstein, R. L. Davies, A. Dressler, R. J. Terlevich, and G. Wegner, Astrophysical

Journal, **326**, 19 (1988).

[21] M. Mevius, S. Buitink, H. Falcke, J. Hörandel, C. W. James, R. McFadden, O. Scholten, K. Singh, B. Stappers, and S. Ter Veen, Nuclear Instruments and Methods in Physics Research A, **662**, 26 (2012).

[22] S. Adrián-Martínez *et al.*, arXiv.org, **1111**, 3473 (2011).

[23] J. Abadie *et al.*, Physical Review D, **81**, 102001 (2010).

[24] The derivation is similar for the opposite case ($x \leq x_{\mathrm{th}}$) where the detector becomes saturated beyond the threshold value and information regarding the exact value of the statistic is lost.